CHAPTER 9

Risk Management

It is important to see distant things as if they were close and to take a distanced view of close things.

—Miyamoto Musashi, *The Book of Five Rings* (1645 CE)

Many of the problems of trading reduce to risk management. In most styles of trading, the trader's job is essentially to manage the risk in trades, focusing on exiting losing trades at the correct points, and letting the upside take care of itself. Risk management is critical, as a few outsized losses can offset the profits from many winning trades; it does not take many errors to completely erode a trading edge. This chapter considers risk management from both practical and theoretical perspectives. We begin with practical, applied tools for risk management and position sizing that traders in all time frames will find useful. Many large losses come from having trades with inappropriate position sizes, and many traders do not understand the impact of trading size on the bottom line. The chapter moves on to a higher-level perspective on risk, incorporating some ideas from modern academic thinking and focusing on a few specific measures of risk, before concluding with a look at some of the less common risks that self-directed traders often overlook.

RISK AND POSITION SIZING

From a practical perspective, there are three main questions to answer with respect to risk. The first two have been addressed in the previous chapter, but the third is supremely important:

1. Where do we place our initial exit orders, both for profits and losses?
2. How do we adjust the trade as it develops through time?
3. How many shares, contracts, or other units do we trade on each position?

First, Know Your Risk

There are few absolutes in trading. Most trading rules are flexible, and many master traders have a rule that basically says, "Know when to break the other rules." However, there is one rule that cannot be broken—it is perhaps the single most important rule in discretionary trading: *Always know where you are getting out of the trade if you are wrong, before you get in*. The choice of exactly where to place the stop will depend on the pattern, the trader, the profit target, the time frame, the specific market, and perhaps a number of other conditions, but the most important thing is that this level is defined at the time of entry.

As you consider your options, it is important to make sure that your stop is placed at a meaningful level that is outside of the range of the market's noise; with very few exceptions initial stops must not be placed too close to the market. A rough guideline is given by the average range of a single bar on the time frame you are trading. If you try to place your initial stops closer than one average bar's range, you are probably working within the noise level and have significantly impaired whatever edge you might have had.

%R and Position Sizing

Once this initial price movement risk level is known, the question of trading size must also be addressed. There is a large body of literature that deals with the theoretical ideas behind asset allocation and position sizing for individual traders. For the most part we will avoid these discussions and will simply limit discussion here to two points: some practical guidelines I have used with success in my own trading and an illustration of why consistent sizing is important.

Many traders are familiar with the *Kelly criterion*, which is a formula that gives the optimal amount to bet in a game of chance, assuming that a number of very important simplifying assumptions hold true. If these assumptions are met, then the Kelly formula will outperform all other approaches, usually dramatically.

> If you are interested in the theoretical background of asset allocation and position sizing for individual traders, good places to start are J. L. Kelly Jr.'s seminal 1956 paper "A New Interpretation of Information Rate" and Ralph Vince's *The Leverage Space Trading Model* (2009).

However (and for actual trading this is very important), if Kelly is applied in cases in which the simplifying assumptions do not hold, then it will all end in tears. For reference, the classic Kelly criterion gives f, the percentage of the account to risk on every trade, by this formula:

$$f = \frac{Odds \times Prob_{win} - Prob_{loss}}{Odds}$$

$$Odds = \frac{Size_{win}}{Size_{loss}}$$

The Kelly criterion and Vince's extension of it, called *optimal f*, are optimized approaches. In both cases, the objective is to maximize the geometric growth rate of the account while minimizing the risk of ruin, or a catastrophic loss of equity. Most optimized approaches like this are extremely aggressive, and large drawdowns must be accepted as a matter of course. In addition, and much more seriously, the theoretical assumptions behind these models are important and they rarely hold up in short-term trading. For instance, most optimized methods assume that each trade is independent of any other trade, but many trading systems experience strings of wins and losses while the market is locked in one regime (trending or trading range), and trade outcomes may show some degree of serial dependence. Furthermore, many of these optimized approaches require inputs like the largest trading loss, which must be based on historical data, and they are all dependent on the assumption that future results will look like the past. If you experience a larger loss in the future and you are using an aggressive, optimized position-sizing methodology, you could be in trouble. If you are going to use any of these approaches in actual trading, make sure you understand the assumptions and the risks involved in violating any of those assumptions.

Many traders find that the risks under an optimized position-sizing scheme are too great to bear. Another possibility could be to set aside a fixed percentage of your trading capital, and apply aggressive sizing rules to that portion while trading the remainder more conservatively. If you decide to do this, you should carefully consider whether the incremental gains offset the additional effort required to carry out this plan.

Fixed Fractional Approaches My approach to position sizing is simple and robust. It is not an optimized method, nor is it meant to be. Rather, it is designed to do a few specific things:

- Define the risk from losing trades.
- Limit the risk from a much larger-than-expected losing trade.
- Limit the risk from a series of losing trades.
- Limit the total risk from a set of highly correlated positions.
- Limit the total amount of equity placed at risk at any one time.
- Allow easy scaling as the account balance changes.

Notice that the focus of this process is on limiting risk, not maximizing returns. This is how you stay in the game. Professional traders know that, if you are trading with an edge, the most important thing is to manage the downside so that no one risk can take you out of business. Profits, to a great extent, will take care of themselves; the first and most important job of any trader is to correctly manage the risk.

The rules are very simple: the risk for every trade is set at a consistent percentage of the account equity. I consider anything under 1 percent to be very conservative, and anything 3 percent or over to be extremely aggressive. When you are thinking about this risk number, it is important to consider the impact of a string of four or five losing trades in a row, or of a single loss five times your anticipated maximum loss. If you are risking

3 percent on a trade and have a disastrous situation where you have a 5× loss, you just lost 15 percent of your equity. In reality, a loss this much larger than expected should be extremely rare, but, even in this extreme case, the account would not be destroyed. If, however, you had been risking 10 percent of the account, you would be down 50 percent and in serious trouble.

Drawdown usually refers to the amount an equity curve, whether for a trader, system, or fund, has retreated from its peak. Drawdowns are a fact of life; there is always some natural fluctuation in an equity curve. Later in this chapter, we examine many of the traditional measures of risk and find them wanting on one level or another. It is easy to make an argument that drawdown is one of the best and one of the truest measures of risk, but, like many other elements of trading performance, future drawdowns must be extrapolated from history. If something changes in the future, the disconnect between historical performance and walk-forward projection can cause traders and managers to dramatically underestimate risk.

Recovering from Drawdowns

One of the mathematical realities traders face is the asymmetry between drawdowns and the percentage returns needed to recover. For some people, this is not always intuitive at first glance—for any given percent drawdown, it takes a much larger percentage return to recover. For any drawdown of D%, the return needed to bring the account back to breakeven may be calculated from this formula:

$$\text{Return needed for recovery} = \frac{D\%}{(1 - D\%)}$$

For small drawdowns, the return needed to recover is only very slightly larger. For instance, a 5 percent drawdown requires a 5.3 percent return to recover, but a trader in a 20 percent drawdown needs to make 25 percent just to break even. For larger drawdowns, the hole gets deeper very quickly—you would have to double your money, with a 100 percent return, to get back to even from a 50 percent drawdown. To put these numbers in perspective, very few traders make 25 percent a year with any regularity, and the chances of doubling your money, without excessive leverage and risks, are very, very small. An effective risk management strategy will focus on minimizing drawdowns, while acknowledging them as a natural and normal part of trading.

Another advantage of fixed fraction approaches is that the actual amount at risk grows or shrinks with the account size. A series of five losses of 4 percent each is not a 20 percent loss to the account because the account was shrinking after each loss. In this case, the actual net loss was 18.5 percent, slightly smaller. The *compound loss* resulting from a string of t losses of N% each is:

$$\textit{Compound loss} = (1 - N\%)^t$$

Note also that this formula would be valid only for *consecutive* losses in an account trading *one position* at a time. If three positions are put on at the same time, they will all be risking N% of the same account balance, so a loss on all three of those *will* be equal to a single 3× loss.

Calculating Trading Size Once you have defined the percentage of account equity to put at risk on any one trade, the next step is to target that risk by having the correct position size (shares, contracts, or dollars invested) on each trade. Most trading patterns or systems will have more or less clearly defined stop points. If you want to risk more or less on the trade, you *cannot* accomplish this by moving the stop further or closer to the market. Essentially, you do not get to choose where to put the stop—that is dictated by the system, but you *do* get to choose how much money to lose by changing your trading size. An example will help to clarify this:

- Assume you are trading a $100,000 account and you want to risk 1 percent on each trade.
- Assume you find a trade to go long a $50.00 stock with a stop loss at $47.50. The difference between the entry and the initial stop is the *per-unit risk* that the trade requires: per share, per contract, per dollar, and so on for each respective market.
- The question is how many shares to trade so that a loss at the market price of $47.50 will equal a loss of 1 percent of the $100,000 account.

$$\text{Trade size} = \text{Desired dollar risk} \div \text{Per-unit risk}$$

In this case, we want to risk 1 percent of $100,000, or $1,000 on each trade, so $1,000 ÷ $2.50 = 400 shares.

Thinking in R multiples This initial risk is very important because it is the base for the *R multiple* for the entire trade. Many traders will find it productive to think of all P&L generated by the trade in terms of this initial R multiple. For instance, if the trade loses the amount that was originally risked, this is a –1R loss, or a loss equal to the initial risk. If the trade has a profit equal to half of this initial risk, we could say that it is a 0.5R trade, and so on. Thinking in R multiples is a valuable skill because it removes much of the pressure of thinking about actual money. Many developing traders find that they relate their initial wins and losses to real-life money situations, as in "I just lost enough on that one trade to make my car payment for three months." Experienced traders are more likely to think of the money as an abstraction or, in some cases, almost more of a score-keeping system. This is a subtle but critical psychological adjustment.

Traders also sometimes have trouble increasing their trading size because of the growing nominal risk. This is a problem for the individual self-directed trader, but can also be a problem on the institutional scale when new client funds come in. There is a big difference psychologically between risking $10,000 on a position and $100,000, but

if they both represent the same percentage of different capital pools and are both the outcome of a 1R trade, it is possible to approach them with a degree of equanimity.

The Effect of Position Sizing

Evaluating the effectiveness of a position-sizing strategy is difficult because doing so requires a good understanding of probabilities in highly path-dependent situations. It is extremely difficult to build intuition in these scenarios, and it is very easy to make mistakes—we are not naturally wired to think about issues like this. Monte Carlo methods provide a useful framework for evaluating different strategies. Let's do a simplified analysis to see what the effect of different position-sizing strategies would be on the same trading system. Assume that you are trading a strategy that wins or loses with equal probability (probability of win = probability of loss = 0.5). When you win, you win 1.2 times the amount risked, and when you lose, you lose 1 times the amount risked; your winners are always exactly +1.2R and your losing trades are always exactly −1.0R. The first question is: would you trade this system? A quick check shows a positive expectancy of 0.1 per trade [0.5(1.2) − 0.5(1.0) = 0.1], so the answer is yes.

Let's further assume that you will trade this system for 250 consecutive trades, at which point you will stop and evaluate your results. It is important to remember that your actual P&L was just one realization of a nearly limitless set of possibilities. This is one reason that we have trouble thinking in probabilities: probabilities are meaningful over a large number of trials, but, in actual practice, we are usually dealing with the realization of one specific outcome. It is easy to focus on that one result and not to realize that it is one of many possibilities and, in some cases, may even be a highly improbable outcome. Our behavior and decisions must be governed by the *most probable* outcomes. *Monte Carlo modeling* or *simulation* can be a useful tool to help build intuition about these path-dependent situations. Each path on the tests that follow represents the outcome of one trader executing these 250 trades in a row; we will repeat the test for each of 1,000 traders. A useful mental trick is to imagine that each path is the P&L for one of 1,000 traders working in parallel universes, since it is difficult to imagine one trader trading the system 1,000 times. Last, most Monte Carlo tests are run on much larger numbers (by several powers of 10). Though the small sets used here may not be big enough to assure convergence to theoretical values, they are large enough to illustrate the relevant points.

So you have this system that you know has a positive expectancy, and you will be trading it in a $100,000 trading account. The only question left is how much to risk on each trade; you have to start somewhere, so let's begin by risking $2,000 on each trade. This seems like a reasonable compromise because it is a number that is significant, but is not extremely large relative to the account value. Before we jump into the Monte Carlo results, see if you can answer this question: What do you expect your winnings should be at the end of the run?

We can answer this by using the expected value of the system, which is 0.1 in this case. This means that, for every dollar wagered, we expect to make $0.10 per trade. (Another way to say the same thing is that for every dollar wagered, we expect to end up

with \$1.10 because the original dollar is returned.) Since we are risking \$2,000 per trade for 250 trades, the relevant math is:

$$\begin{aligned} \textit{Ending P\&L} &= \textit{Expected value} \times \textit{risk per trade} \times \textit{number of trades} \\ &= 0.10 \times \$2{,}000 \times 250 \\ &= \$50{,}000 \end{aligned}$$

So, *on average*, we can expect to end up with \$150,000 in the account (\$100,000 starting value + \$50,000 profit) at the end of 250 trades.

For each of the 1,000 runs, the final, ending P&L was recorded, and Table 9.1 presents summary statistics for these ending values. Most of the statistics in the table are self-explanatory. With one exception, these statistics are blind to what happens in the middle of the run. If a trader was up \$5,000,000 at one point but the run ended up only \$5,000, the only number evaluated is the ending \$5,000.

The exception is that an account cannot continue to trade if it loses all its funds at any point in the trial, so each run is monitored for this condition. If, at any point in the series of 250 trades, a run reaches a zero or negative balance, it will cease trading and that value will be presented as the terminal value for that run. The number of the 1,000 traders who go bankrupt is recorded, and the percentage is presented at the bottom of the table.

$$\text{Percent of accounts bankrupt} = \text{Number of accounts with terminal balance} \leq \$0/1{,}000$$

We also will record the number of accounts than end up with a greater than 75 percent loss from the starting value. Though this is an arbitrary cutoff, it will serve to measure some of the downside risks in some of the more extreme scenarios later on. In reality, an account is in trouble long before it reaches a 75 percent drawdown. Consider the results in Table 9.1.

On average, we end up with \$149,259, which is close to the theoretical \$150,000 from the expected value equation. (Keep in mind that 1,000 test runs is actually a small number; had we done 100,000 or 1,000,000 runs, this number would almost certainly have

TABLE 9.1 Results of 1,000 Monte Carlo Runs Risking \$2,000 per Trade in \$100,000 Account

Mean Terminal Value	149,259
Median Terminal Value	152,000
Standard Deviation of Terminal Value	33,760
Coefficient of Variation	0.23
Mean of Maximum Value	159,070
Mean of Minimum Value	89,254
Highest Terminal Value	253,200
Lowest Terminal Value	50,800
Percent Terminal Drawdown ≥75%	0.0%
Percent of Accounts Bankrupt	0.0%

been closer to $150,000.) The standard deviation of terminal value tells us that 68 percent of the time, the terminal value should fall within +/−$33,760 of $149,259 *if* the terminal values are normally distributed. (Though not necessary in the context of Monte Carlo modeling, it pays to check assumptions of normality as a matter of course. In this case, they are, in fact, normally distributed.) The *coefficient of variation* is not particularly meaningful in this one case, but it gives us a quick tool to make risk-adjusted comparisons across different scenarios.

Make sure you understand the mean of maximum and mean of minimum terminal values: the average of the highest and lowest points reached by each of the 1,000 acounts over 250 independent trades risking $2,000 per trade. These are averages, and there will certainly be surprising deviations, especially in a larger number of trials. For instance, in this test we had one run in which the account balance reached $12,400 at one point. (This is not visible in the table because the account recovered significantly by the end of that specific trial.) At this point you should be saying something like, "What? I thought we were trading something that has a positive expectancy." Yes, this is a simple system with a clear positive expectancy, but one trader trading it would have experienced an 87.5 percent drawdown, purely due to the slings and arrows of outrageous fortune! This is an important reality check, and it gives us some sense of how much variation can exist within a positive expectancy framework. In this case, the account recovered to close at $55,200, but that also might not have happened in another universe. (Note that this is also not the lowest closing value of all accounts.) Is it possible that you can trade a system with a positive expectancy, make no mistakes trading it, and still go bankrupt? Absolutely. If so, how do we really evaluate our trading performance and separate skill from luck? This is an important question to think about—though there are no certain answers, Chapter 12 will offer some guidelines and helpful measures.

In Table 9.1, the highest and lowest terminal values are exactly what they say: the largest and smallest account values at the end of the 250-trade test; there will almost certainly have been significantly higher and lower values at other points in the run. (We already know that one of these accounts traded down to $12,400 before recovering to close much higher, and one of them traded up to $265,200 at one point during the test. (Neither of those values is shown in the table.)) Last, we see that none of the accounts ended in a greater than 75 percent drawdown, though we have no idea from this table how many may have crossed that barrier during their life cycles, and we see that none of the accounts went bankrupt (i.e., reached a zero or negative balance at any point during the trading, at which point trading would have been terminated).

So what conclusions can we draw from this? Well, first of all, math works—the theoretical expected value calculation gave a number very close to the mean of the terminal values, and the difference is easily explained by normal variation. However, this concept of expected value fails to capture the degree of variation possible within a positive expectancy framework. Few readers would expect to trade a system with a verifiable, valid trading edge and to lose nearly 90 percent of their accounts to random bad luck, but this can and does happen. There are also offsetting happy surprises to the upside, but it is very hard to develop an intuitive grasp of the contribution of randomness to our

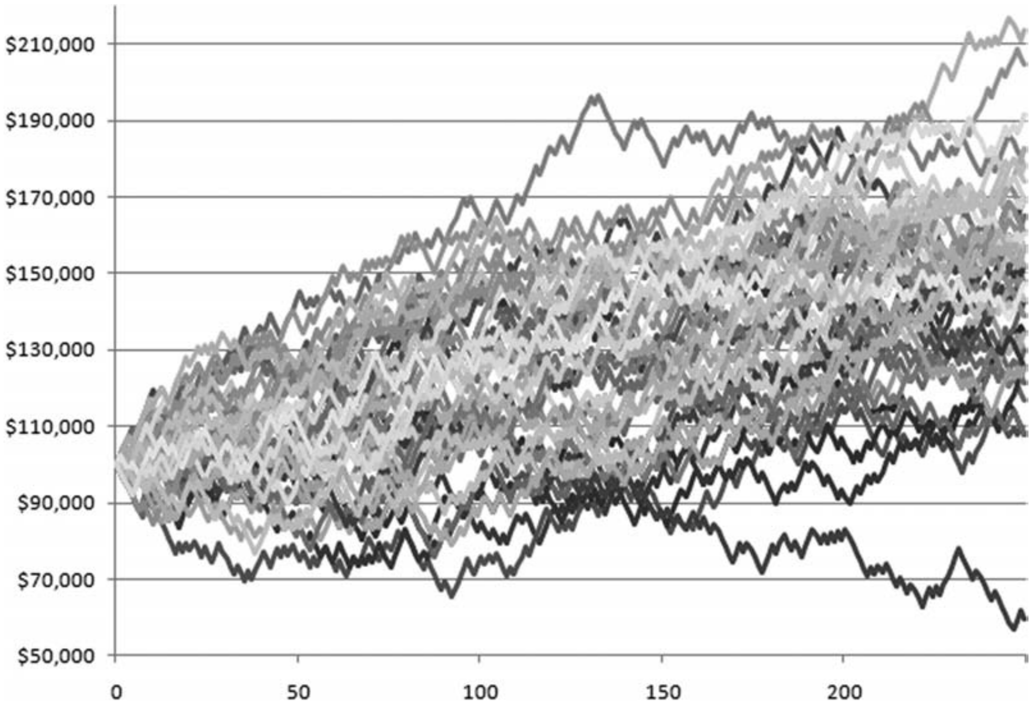

FIGURE 9.1 Fifty Sample Runs through the Monte Carlo Tree

trading results. Figure 9.1 shows the paths of the first 50 runs through this test. Though it is difficult to read such a crowded graph, notice that most of the paths do cluster around a rising central value, as would be expected. However, also notice that there are upside and downside outliers, some of which are quite extreme.

So, with that background, let's consider what would happen if we reran the test, changing only the amount risked on each trade. Table 9.2 shows results for several

TABLE 9.2 Monte Carlo Risking $x per Trade

	x = 2,000	x = 3,000	x = 4,000	x = 5,000
Mean Terminal Value	149,259	173,854	198,286	221,894
Median Terminal Value	152,000	178,000	204,000	230,000
Std Dev of Terminal Value	33,760	50,746	68,107	86,911
Coefficient of Variation	0.23	0.29	0.34	0.39
Mean of Max Value	159,070	188,604	218,139	247,448
Mean of Min Value	89,254	83,912	78,650	73,638
Highest Terminal Value	253,200	329,800	406,400	483,000
Lowest Terminal Value	50,800	(1,400)	(3,200)	(4,000)
Percent Terminal DD ≥75%	0.0%	0.1%	0.8%	2.3%
Percent of Accounts Bankrupt	0.0%	0.1%	0.7%	2.2%
Terminal Value − $E()$	(741)	(1,146)	(1,714)	(3,106)

different per-trade risks ($x =$). Note that the *only* thing that changed is this risk level; the trades for these runs were generated with a pseudo-random number generator, and the same random sequence was used in each set of 1,000 trades. The win/loss sequences are the same for the set of 1,000 trials and for each column in Table 9.2. (On a technical note, this kind of reproducibility is important in Monte Carlo testing, and is a strong argument against building these tests using Excel's built-in random number generator.)

What do we notice here? First, the simplistic understanding that some traders have that "the more you risk, the more you make" does seem to have some validity. As we increase the amount risked per trade, the average terminal value increases, as does the mean of the maximum value; risking $5,000 per trade gives us almost 50 percent more profit than our initial $2,000 risk. However, there is a cost. The standard deviation of terminal values increases faster than the mean, which can be seen by the steadily risking coefficient of variation. If we accept, for a moment, the standard deviation as a measure of risk, we are taking on additional units of risk and not being adequately compensated by higher returns. At the $5,000 risk level, we have a disturbing number of accounts that went bankrupt.

A new line has been added to the table that shows the difference between the terminal value and the expected value. The difference was very small at the $2,000 risk level, but it increases with the risk level. One reason for this departure is that the bankruptcy limit makes the test, and actual trading, asymmetrical. If an account hits this barrier, it is removed from the test and that value becomes its terminal value; this is also the reality of capital constraints in trading and money management. Otherwise, martingale betting strategies (where you double your bet size each time you take a loss) would work, but in reality, traders using this type of strategy invariably go broke if they trade long enough.

Emboldened by the observation that more risk equals more profits, let's say you decide to, as we kindly say in the vernacular, "go nuts." You increase risk to the reckless levels in Table 9.3, and finally some truths become very clear. At some point, with increasing risk, the party is over—the rising losses and risks of bankruptcy start to outpace

TABLE 9.3 Aggressive Risk Levels in the Monte Carlo

	$x = 6{,}000$	$x = 8{,}000$	$x = 10{,}000$	$x = 25{,}000$
Mean Terminal Value	243,140	281,387	314,822	463,980
Median Terminal Value	242,800	290,400	338,000	310,000
Std Dev of Terminal Value	108,326	153,502	200,806	519,633
Coefficient of Variation	0.45	0.55	0.64	1.12
Mean of Max Value	275,360	327,832	376,324	610,805
Mean of Min Value	68,993	61,058	54,662	28,480
Highest Terminal Value	559,600	712,800	866,000	2,015,000
Lowest Terminal Value	(5,600)	(7,200)	(8,000)	(20,000)
Percent Terminal DD $\geq$75%	5.2%	11.0%	17.6%	47.7%
Percent of Accounts Bankrupt	5.0%	11.0%	17.6%	47.7%
Terminal Value $-$ $E()$	(6,860)	(18,613)	(35,178)	(261,020)

any incremental gains. Somewhere between $8,000 and $10,000 risk per trade, we finally reach a mean terminal value that is double what we have with our modest $2,000 trading risk, but we also have somewhere around 15 percent of the accounts going bankrupt. Though decisions have to be made within your own risk tolerance, it is hard to imagine this being acceptable. Consider the $25,000 per trade risk as a very extreme example. For one thing, you are risking a quarter of your initial account balance on each trade, which means you can be wrong three times before being taken out of the game on the fourth bad trade. What are the chances of that happening? It turns out the chances are pretty good, as nearly half of the accounts go bankrupt at this excessive risk level. Also, we are now risking 11.5 times our initial $2,000 risk, and the mean terminal value is not really *that* much higher—certainly not enough to justify the excessive risk of going bankrupt.

Consider one more thing about this $25,000 scenario: there was at least one account in this test that ran the starting capital of $100,000 up to over $2,000,000 (not visible in the summary table), at the same time that half the other accounts were going bankrupt. It is always possible to find one exceptional example.

Fixed Fractional Position Sizing in Action Though risking a fixed amount on each trade is an improvement over having no system at all, there are some problems with this plan. For instance, if the account balance shrinks or grows dramatically, the initial dollar amount risked may lose some relevance. If your account doubles, why are you still risking the same amount per trade?

Fixed fractional risk is one solution to this problem: always risking a fixed percentage of the trading account's equity on each trade. Let's begin this investigation with a 2 percent risk, which, at the beginning of the run, will be the equivalent of the $2,000 risk in Table 9.1. Keep in mind that the amount risked will change on every trade, as the account balance waxes and wanes. As we have winners, the account gets bigger, so we are risking more on the next trade; conversely, as we run into a string of losers and the account shrinks, the risk on each trade will also shrink. Table 9.4 shows the results of this 2

TABLE 9.4 Fixed Fractional and Fixed Position Sizes Compared

	Fixed 2,000	2% of Equity
Mean Terminal Value	149,259	163,033
Median Terminal Value	152,000	158,317
Std Dev of Terminal Value	33,760	56,266
Coefficient of Variation	0.23	0.35
Mean of Max Value	159,070	179,236
Mean of Min Value	89,254	89,385
Highest Terminal Value	253,200	434,738
Lowest Terminal Value	50,800	57,654
Percent Terminal DD ≥75%	0.0%	0.0%
Percent of Accounts Bankrupt	0.0%	0.0%

percent fixed fractional sizing, and, though they are not directly comparable, reproduces the numbers from Table 9.1 for comparison.

A few things jump out here. First of all, these numbers do seem to be roughly comparable at first glance. We see that both the mean and the median values have increased, which is to be expected because, as the account grows—which, on average, it does with this positive expectancy system—fixed fractional sizing allows you to take on more risk. This is essentially a way to leverage your winners and to use accumulated trading profits to fund further risk. We note that the standard deviation has increased fairly dramatically; the coefficient of variation does not look good. In fact, if they were directly comparable, Table 9.2 suggests we should be able to hit a mean terminal value of around \$200,000 for a coefficient of variation of 0.35 with a fixed dollar amount risk plan. Perhaps we are a little disappointed with the fixed fractional results, which seem to have increased volatility of returns without compensating for the extra risk. In fact, the fixed fractional approach would have a lower Sharpe ratio for a commensurate return. Naive rules would suggest that this is lower risk-adjusted return, but hold off on making any decisions for a moment. There just may be more to this story.

Digging a little deeper, we see that the mean of maximum value has increased fairly significantly, without any real change in the mean of minimum value. In fact, both the mean of minimum value and the lowest terminal value from the run have *increased*, and we also see that the highest terminal value is about 1.7 times what it was with the fixed sizing. Figure 9.2 gives some insight into what is going on.

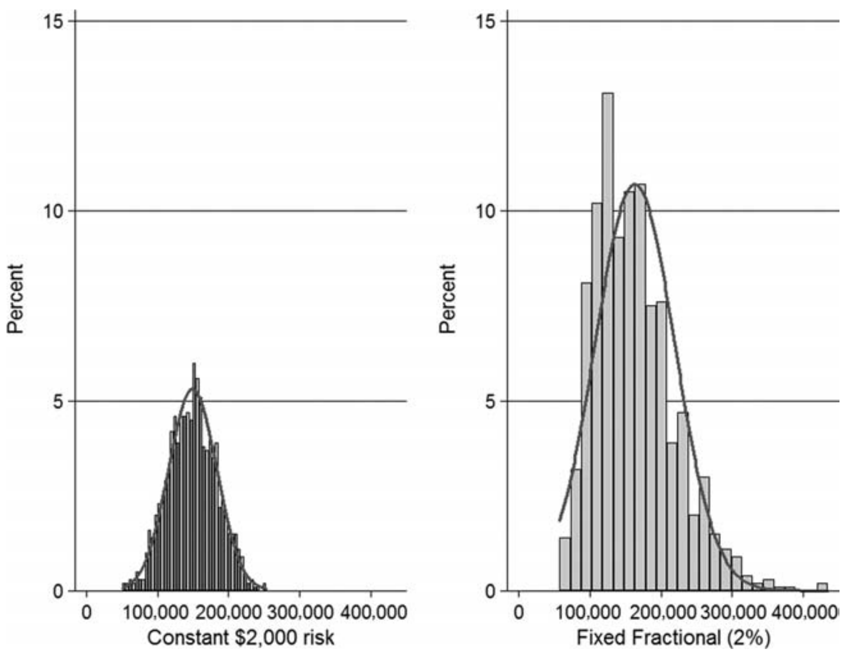

FIGURE 9.2 Distribution of Terminal Values for Fixed-Dollar and Fixed-Fractional Position-Sizing Plans

In Figure 9.2, the histograms also include a normal distribution curve for comparison with the returns. We established earlier that the constant risk scenario produced ending values that did seem to be normally distributed. However, even a casual glance at the fixed fractional distribution shows that it almost certainly is not normally distributed. (In this case, the returns have a skew of 0.95 and kurtosis of 4.3. The Shapiro-Wilk test confirms with a z value of 8.5, giving strong evidence of nonnormality. The returns of the fixed fractional plan are *lognormally* distributed.) The most obvious point about the fixed fractional distribution is that it is no longer symmetrical; the variation is concentrated in a long right (positive) tail. *This is critically important:* the asymmetrical risk profile compromises the relevance of simplistic measures such as the Sharpe ratio or the coefficient of variation. In this case, the increased risk from the fixed fractional approach is a good thing; nearly all of the extra variability is potential upside.

You Can't Go Broke? Really? Advocates of fixed fractional approaches often point out that it is mathematically impossible to take an account to zero using these approaches. As the account balance shrinks, you are risking an ever-smaller percentage of the declining balance. This is slightly reminiscent of Zeno's paradox in which you walk halfway to a fixed point, then half of that distance, then half again. Repeated into infinity, it seems as though you would never reach your goal. The people who say that fixed fractional approaches prevent the account from going to zero are absolutely correct (in the absence of commissions and other frictions), but, unfortunately, it is a completely meaningless argument. If we did 250 losing trades in a row, risking 5 percent of our equity on every single trade, our $100,000 trading account would not be completely gone—we would, in fact, have about $0.28 left. This is comforting if you consider a 99.9997 percent loss to be significantly better than a 100 percent loss, but the distinction is more than a little pedantic.

The advantage of a fixed fractional approach is expanded upside and dampened downside potential, not that it protects you from going bankrupt per se. This is also the reason for the pragmatic 75 percent ruined account barrier set in all of the Monte Carlo tests. In practice, most traders and investors will be uncomfortable with much smaller drawdowns, but this is a useful reference point across these tests.

Using Different Fixed Fractions We used a 2 percent fixed fractional size as a starting point without any real analysis or consideration; it is possible that other percentages could give better results. Table 9.5 shows the Monte Carlo results for several different fixed fractional sizes, including some that are a little crazy. When we were risking a constant $25,000 on each trade, nearly half of our accounts closed below $25,000, the average closing value was $463,980, and some lucky trader made over $2,000,000 at the end of the run. If we compare that to the fixed fractional 25 percent (again, not a direct comparison because that number changes with the account balance), we see the power of the adaptive bet size. Now, 66.5 percent of the accounts close with a greater than $75,000 loss, which is not good. (It's actually a little worse than you might think from the table. Only 281 of the 1,000 accounts actually made money. Over 450 of

TABLE 9.5 Several Fixed Fractional Monte Carlos

	4%	8.333%	10%	12%	25%
Mean Terminal Value	263,545	711,375	1,014,163	1,508,351	5,280,790
Median Terminal Value	222,225	307,825	300,832	262,485	5,630
Std Dev of Terminal Value	193,846	1,431,904	2,821,550	6,005,588	79,775,366
Coefficient of Variation	0.74	2.01	2.78	3.98	15.11
Mean of Max Value	316,875	1,046,380	1,636,924	2,781,394	62,017,391
Mean of Min Value	79,046	58,366	51,237	43,292	9,444
Highest Terminal Value	1,670,569	20,392,572	46,000,315	109,597,717	1,757,129,394
Lowest Terminal Value	29,561	4,647	1,967	629	0
Percent Terminal DD ≥75%	0.0%	4.6%	8.0%	15.6%	66.5%

the initial 1,000 accounts closed with less than $2,000—certainly a catastrophic loss by any reckoning.) However, the average closing value is well over $5,000,000, and one of the accounts made almost two *billion* dollars. This is not a bad run from a starting value of $100,000.

The 8.333 percent example in Table 9.5 is a special case: it is the Kelly number for this particular system. It is worth pointing out that this is exactly the kind of situation in which the Kelly criterion can be safely applied. The wins and losses are all the same size, and there are never any surprises in the form of larger-than-expected wins or losses. In addition, each trade is completely independent of preceding trades. In other words, there are never longer strings of winners and losers than you would expect to see in random data. The Kelly number produces some pretty impressive results, with mean and median terminal values far exceeding anything we were able to hit with the (reasonable) constant risk scenarios. In addition, the Kelly criterion opens the door to truly outstanding performance; one of our Kelly-sized accounts would have run the $100,000 initial investment up to over $20 million by the end of the run. As you might expect from an aggressive strategy, there is also significant risk of loss, as 4.6 percent of the accounts closed under $25,000. Even so, there may be situations in which a trader could accept the risks of an optimized strategy in return for the potential reward. This might depend on the trader's risk tolerance and investment goals for the account. For instance, is the trading account your entire net worth, or is it just a small subaccount that you intend to press aggressively? These questions go beyond the math behind the strategies, but they are important questions to consider before you actually put capital at risk in the market.

As you might imagine, since it is an optimized, aggressive number, bad things start to happen when we risk more than the Kelly criterion suggests. Though the general

principle of "risk more, make more" is *technically* true, as you can see from the mean terminal values for the 10, 12, and 25 percent risk levels in the table, the declining medians tell a different story. We cannot see it in the table, but the distribution of these ending balances becomes unacceptable, and the payoff of our well-considered positive expectance trading strategy starts to look more like a lottery ticket as we risk a higher percentage of the account. Do not be misled by the sheer size of the winners, because, as you know by now, the extremely low probabilities associated with those outcomes moderate the expected value.

There is one last important thing to consider. Many traders like to change their bet size based on any number of factors. This can be a well-thought-out and disciplined element of a trading strategy or of trading certain markets. For instance, it might make sense to trade certain kinds of patterns on smaller risk (for instance, the failure tests in Chapter 6), and some traders might want to approach illiquid markets with smaller risk. Too many discretionary traders make emotional decisions about risk without any real analysis, varying their risk based on their impressions of how good a trade is likely to be. If you are making emotional decisions about risk, you are almost certainly making suboptimal decisions. If you are going to intervene and meddle with position sizing and risk, it is important that you do two things: One, understand the impact of random sizing on a system and a strategy. Two, keep careful records and do objective analysis to be sure that your actions are actually adding something of value compared to the unadjusted position sizing rules.

Table 9.6 shows the effect of randomly changing the bet size on each trade. The 4 percent column is reproduced from Table 9.5 for comparison, and then three other tests were run where each bet size was a random value between 0 and 8 percent (~Uniform [0, 8%]). On average, the random bets were still 4 percent, but they varied between 0 percent, where the trader skipped the trade altogether, to 8 percent, where the trader sized up and took on double risk. Notice that in all cases, the randomly sized bets underperformed the simple fixed fractional 4 percent. Of course, it does not have to be that way. The

TABLE 9.6 Examples of the Effect of Randomly Varying Bet Sizes

	4% Nonrandom	**Rand1 [0%–8%]**	**Rand2 [0%–8%]**	**Rand3 [0%–8%]**
Mean Terminal Value	263,545	262,190	257,249	254,106
Median Terminal Value	222,225	192,474	198,389	187,049
Std Dev of Terminal Value	193,846	239,548	211,231	220,697
Coefficient of Variation	0.74	0.91	0.82	0.87
Mean of Max Value	316,875	333,660	326,773	324,438
Mean of Min Value	79,046	74,116	74,063	73,836
Highest Terminal Value	1,670,569	2,062,297	1,886,388	2,001,251
Lowest Terminal Value	29,561	14,833	18,562	20,599
Percent Terminal DD ≥75%	0.0%	0.4%	0.2%	0.3%
Percent of Accounts Bankrupt	0.0%	0.0%	0.0%	0.0%

random sizer could have gotten lucky, but he *usually* will not be lucky, which is what matters. Note that in nearly all scenarios where the mean of the randomly sized scenarios is not much different from the nonrandom baseline, only the variability increases and it does so symmetrically, meaning that, in this case, standard deviations actually are a good proxy for risk. Changing bet sizes simply introduces another degree of freedom into the equation and brings more randomness to the bottom line. Most traders will find that their interventions in position sizing are actually harming their performance because they are making the wrong decisions at the wrong time.

Other Approaches to Position Sizing There are other approaches and ideas for position sizing, which may be more appropriate for some situations than others.

- *Fixed-percentage allocation:* In portfolio construction and optimization, the values of various assets and asset classes are considered as a percentage of the overall portfolio. There are well-established formulas for understanding the theoretical contribution of each asset to the overall portfolio in terms of volatility, returns, and correlation. For the active trader, I would argue that this number is fairly meaningless. Depending on the risk profile of the market and the specific position, a position that is 5 percent of the portfolio might carry more risk than one that is 50 percent. This is rarely a meaningful measure of the risk associated with shorter-term, technically motivated positions.
- *Margin as a fixed percentage of account:* This is the futures trader's equivalent of the equity trader's percentage of equity allocation in a portfolio. Positions are sized so that the margin required for those positions is a consistent percentage of the account. Exchanges set margins based off their measures and perceptions of risk associated with each market, so, in theory, there could be some justification for this approach. In practice, though, it is a disaster waiting to happen. Volatility is a consideration in the margin calculation, but many other factors determine margins. There are much better ways to size positions.
- *Equivalent volatility strategies:* This strategy has received a lot of attention from portfolio managers and allocators over the past several years. Variations of this concept are sometimes called "equal risk" strategies, and it is also well known that Richard Dennis's and William Eckhardt's Turtles used a variant of this sizing strategy. Basically, you calculate a volatility measure for each market—Average True Range (ATR) is a good standard—and then size each position so that its daily contribution to the portfolio will be the same. You will trade a much larger size in a quiet stock or commodity and a much smaller size in a more volatile one. The point of this sizing is to equalize the impact of each market on the daily portfolio P&L, and it does that very well. It must be emphasized that this strategy does not truly measure and equalize the risk in each position, nor does it consider cases where volatility might be compressed and the market is overdue for a range expansion move. In the wrong hands, this kind of sizing is dangerous, as traders will be putting on very large size at points where the market is highly compressed.

There are many possibilities and choices to be made in position sizing. For the self-directed discretionary trader, this is an important element of the trading plan that deserves careful consideration and study.

THEORETICAL PERSPECTIVES ON RISK

Mathematically, risk is part of the expected value function, and can be defined as:

$$\text{Risk} = \text{Probability of loss} \times \text{Expected size of loss}$$

Though this is a simple equation, there is a profound truth hidden here: the key to properly understanding risk is that the magnitude of the risk depends on both the *probability* and the *size* of the loss (see Figure 9.3). Risks that are very rare and also carry no serious consequences are usually insignificant and can be ignored. However, make sure the consequences are really as small as you think they are; it is easy to miscalculate because these events are infrequent. In addition, consequences are not always consistent—is it possible that there is a small subset of these events that carry more serious outcomes? Risks that are common with low consequences demand careful attention and scrutiny. Traders tend to ignore these risks because they are clearly defined and seem small enough to be insignificant, but a large number of these can eat away at a trading account over time. Examples of these kinds of risks in trading might be normal slippage and transaction costs.

High-probability, high-consequence risks are also usually easy to avoid. These are risks that most people would characterize as stupid and inviting disaster. Most people will look at a risk like this and simply say, "Why would you do that?" Traders usually do not have issues with these risks because they effectively eliminate traders from the industry through a quick and efficient process of natural selection. If you take those kinds of risks as a trader, you will soon not be a trader.

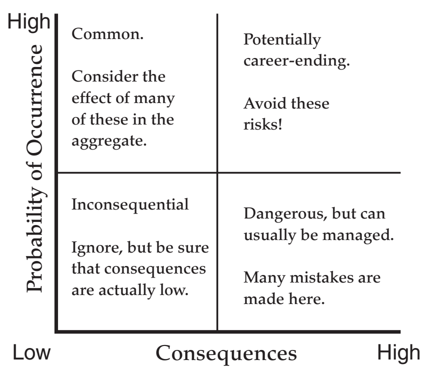

FIGURE 9.3 Probability/Consequence Risk Grid

In general, humans do not have a good internal framework for handling the calculus of risk and probability, especially in high-consequence, low-probability situations. The most dangerous risks for traders lie in this quandarnt. Some people obsess over these risks, but other times these risks are completely ignored. It is easy to mistake extreme improbability for impossibility. Take, for instance, the chance of the planet Earth being struck by an asteroid. It is difficult to imagine a more disastrous event—but the chance of it happening in our lifetime is vanishingly small. Should we be doing more to prepare for and to protect against this risk and at what expense?

I don't know the answer to that question, but I do know that many traders take risks with their portfolios that are, in a financial context, just as severe. A trader selling a lot of naked option premiums could have a single loss that wipes out decades of profits and puts him into personal bankruptcy, but the chance of that happening on any one trading day is very small. This trader probably thinks he can ignore the risk, and he is right, until it happens. Realize that the most serious dangers and risks are these risks that are extremely rare with dire consequences.

Thinking in Sample Sizes

One of the core problems is that thinking in terms of probabilities requires the ability to think in large sample sizes, while we are faced with the concrete realization of a single outcome in the real world. For instance, a common line of thought goes something like this: yes, the probability of that plane crashing might be very small, but that probability doesn't matter to the passengers who happen to be on a plane that *does* crash. This kind of thinking focuses on a single outcome only, not realizing that the true risk depends on both the severity of that outcome and the probability of it occurring. The true risk in this case is the risk present for millions of airline passengers, not those specific, unfortunate passengers who happen to be on a plane with a catastrophic mechanical failure.

It is possible that you play a game in which you have a positive expectancy and lose on several trials in a row, just as it is possible for profitable traders to have losing streaks. If this were your only trial, or if you made only one or two trades, you could end up with a loss due to the luck of the draw. It is counterproductive to focus on any specific outcome; some people say that you can't put probabilities in your bank account, but this is wrong—that is exactly what skilled traders do. Over a large set of trades, traders *do* put probabilities in their bank accounts. The outcome of any one trade, winning or losing, doesn't matter. What does matter is the sum of results over many trades.

Traders make these probability/payoff miscalculations all the time. One of the classic examples is the choice to use very tight stops because the trader wants low-risk trades, not realizing that a tight stop is often a very small but certain loss. Over a large set of trades, the risk of the those many small losses may actually be much higher than the risk associated with a larger stop, because the wide stop will be hit much less often. However, if that large stop happened to be hit on the very first trade, your perception would likely be skewed, and you would probably approach future scenarios with a

different mind-set. The key here is to realize that your intuition and perceptions regarding probabilities are unreliable and must be backed up by logic, reasoning, and testing of relevant scenarios.

Uncertainty as Risk

One commonly used definition of risk in many fields of finance is that risk is uncertainty. For instance, when different assets are being evaluated, the standard deviations of their returns are often used as a proxy measure of risk. This raises a few issues that are often ignored: First of all, is this right? Is uncertainty an appropriate measure of risk? This is what is taught in business schools around the world, so it is often accepted by practitioners without question. True, there are applications in other fields of finance where uncertainty may be a good measure of risk. For instance, in corporate finance, managers often prefer a project that has a lower, more certain return to one that has much higher potential payout but with more uncertainty. Certainty of cash flows is the goal in some cases because it allows for the most efficient allocation of resources—surprises are not good. Though this is appropriate from the manager's perspective, it may not be the best standard to apply to trading and investment management, where strategies may sometimes be completely dependent on positive surprises and outliers.

This desire for certainty finds its way into investment management when portfolio managers use the standard deviation of an asset's returns, measured from past returns, as a predictor of *future* risk. This is potentially flawed on several fronts. First, there is no guarantee that future returns will carry similar variation to historical returns. In fact, it is easy to construct a study comparing standard deviations and distributions for assets across different time windows that will reveal the folly of this assumption. It is also entirely possible that an asset may carry risks that are not visible in the return series. For instance, the returns of a trader who simply writes option premiums every month are probably extremely consistent. A strategy like this may very well show many consecutive positive months (or years) with little or no variation, but there is no accounting for the hidden risk of an eventual catastrophic loss—execute this strategy long enough, and there will be a day of reckoning. Other examples might be complex derivatives, such as those that precipitated the 2007–2009 financial crisis, that have embedded structural risks; a return series is irrelevant for such an asset, especially if the market for them is illiquid and inactive. Quantitatively, there is no way to understand these risks from a simple examination of returns, but qualitative assessment might reveal some of the flaws and hidden risks in many of these cases.

Symmetry of risk is another issue that standard models may not capture very well. In the case of the short premium options trader, there is unseen large left-tail risk, but in other cases, the risks may be dramatically skewed to the upside. Which of these carries more risk: a fund that will return 5 percent annually with a 10 percent standard deviation or a fund that will return 5 percent annually with an asymmetrical risk profile extending 5 percent downward and 15 percent upward? The second fund has dramatically higher

variability, but all of the extra "risk" is skewed to the upside—in this case, variability is opportunity, not risk, as we saw earlier with fixed fractional position sizing. Simplistic measures of risk will consistently misprice asymmetrical risks in the real world.

MISUNDERSTOOD RISK

If the academic world and market researchers are somewhat divided on how to measure risk, it is no wonder that traders are a little confused. One of the key skills of competent trading is excellent risk management; this can be achieved only if the risks are fully understood and accepted. It is important for traders to think deeply about the nature of the risks they are assuming, and to put them in the context of their expected returns.

Generating Positive Expected Value

We can redefine the expected value equation like this:

$$E = (Prob_{win} \times Size_{win}) - (Prob_{loss} \times Size_{loss}) \therefore$$
$$E = Reward - Risk$$
$$\text{Win ratio} = \frac{Prob_{win}}{Prob_{loss}}$$
$$\text{Reward/risk ratio} = \frac{Size_{win}}{Size_{loss}}$$

where the probabilities and sizes of wins and losses are averages over a large set of trades.

The job of a trader seeking to take consistent money out of the market can be simplified to making the E in this equation bigger than zero, or, more formally, to achieving a positive expected value over a large set of trades. It should be clear that there are two paths leading to this same end: the trader can get there through having either a high win ratio or a high reward/risk ratio. Any combination of the two that results in a positive expected value will make money in the absence of transaction costs. There is no inherent advantage to high-probability trading, and there is also no reason to think that "high risk/reward" trades are better.

The "Risk" of Risk/Reward

Mark Douglas (2001) makes a good argument that one of the major reasons traders struggle psychologically is that they are unable to fully accept the risk in any trade. I would go one step further and say that the reason they are unable to accept the risk is that they don't really *understand* that risk. This may be a radical idea, but, over a large

set of trades, I would argue that the "risk" part of the risk/reward equation is not actually true risk.

We live in a world of probabilities, not certainties. No matter how good a trader you are or how good your system is, the outcome of any one trade is more or less a coin flip. You can be a great trader and have three or four losing trades in a row, but this is not *actually a risk.* The losses on those trades are not true risk; they are the cost of doing business. Consider the case of a retail business with a simple business model: the business buys things (inventory), marks them up, and hopefully sells them for a little more than it paid for them. If this were your business, would you consider the money you spent on inventory to be risk, or is it simply a planned expense? Certainly, there *are* risks associated with your business: you can buy insurance to guard against fire, theft, flood, or other damage, but inventory expense is a regular, recurring expense as are normal trading losses.

Some might call this an argument over semantics, but it is actually a very important distinction. When you enter a specific trade, you do not know if you will make or lose money on the trade, so there is a possibility of loss, but over a large enough set of trades, there is a near certainty of a profit *if* you are trading with a positive expectation. This loss is not, in the long run, a risk. (If you are not trading with a positive expectation, then you have a different problem.) Traders are paid for assuming the correct risks at the correct time within a larger-scale, positive expectancy framework. Until you fully assimilate this truth, you will struggle—this is probably the one belief the separates profitable discretionary traders from everyone else.

PRACTICAL RISKS IN TRADING

There are many very real risks in trading. Some of them are obvious, many of them are misunderstood, and some of them are unknowable. The purpose of this section is to get you thinking about these risks, some of which you will encounter every day, and some of which you may meet only once in your career. This is not an exhaustive list of all possible risks, but it does give you some ideas of things to think about.

Trading Without an Edge

One of the worst possible sources of risk must be acknowledged up front—perhaps the trader actually does not have an edge in the market. Larry Harris (2002) calls traders who trade without an edge *futile traders*, and has this to say about them:

> *Futile traders expect to profit from trading, but they do not profit on average. They cannot recognize the difference between their expectations and their results. They may be irrational, they may have poor information about their results . . . or they may be of limited mental capacity.*

He goes on to define *inefficient traders*, saying:

> *Inefficient traders lack the skills, analytic resources, and access to information to trade profitably. They may do everything that profitable traders do, but they do not do it well enough to trade profitably. . . . [Their] profits are not sufficient to cover their losses to more skilled and better-informed traders. Inefficient traders generally make poor decisions about when to trade and when to refrain from trading.*

Without a real trading edge, everything else is a waste of time. Nothing will help; the trader is simply doomed to lose money because the markets are actually worse than our theoretical random walk zero-sum games—the markets are actually a *negative*-sum game due to trading frictions and transaction costs. It is impossible to make money trading without a real edge in the market, and, furthermore, if you don't know what your edge is, you don't have one.

A surprising number of retail traders operate without an edge, and they consistently bleed money. Why would anyone do that? As Harris speculated, perhaps they may be of limited mental capacity, but probably a more common reason is poor record keeping. Few retail traders keep professional-level records tracking every aspect of their performance over a significant period of time. Though record keeping seems like an incredibly mundane topic, the discipline of tracking performance is a core skill of professional trading. It is not uncommon to see working traders devote more attention to the process of record keeping than to actual trading. If you are keeping good records and carefully tracking your trade stats, the numbers will not lie.

Many traders fail and trade without an edge because they do not know how to do the correct analysis to understand their edge, or they make any of a number of mathematic errors in that analysis. (See Chapter 12 for some specific tools and a structured approach to this analysis.) Traders may work hard and sincerely believe they are trading with an edge, but they often are not, and they are confused by the fact that their trading accounts are continually decreasing in value. Many discretionary traders will resist the idea of a systematic, structured approach to trading, but I am not advocating that—I *am*, however, pointing out the necessity of that approach in analysis and in designing your trading system. Unless you do this analysis, how can you be sure you are really trading with an enduring edge in the market? Fail to do this, and you will lose.

Execution Risk

Execution risk is a topic that covers a number of events, including some that are out of the trader's control. Every trader should approach the market with a precise plan for when and where to enter and exit the market on every trade. Execution risk is the divergence between those intended points and the actual prices received. Though these differences will sometimes be in the trader's favor, on balance they will not be. *Slippage*, missed trades, and miscellaneous execution errors add up; collectively, these types of

mistakes can become a trading friction that will significantly erode the trader's edge in the market.

The sources of these errors will vary from trader to trader and from market to market and will even be different at different points in a trader's career. For instance, newer traders often miss trades because they are nervous about the (false?) risk associated with the trade. They may hesitate or they may jump the gun and execute trades when their methodology does not actually call for an entry. Some traders randomly alter sizes, sometimes taking a few trades on much larger size, or they may try multiple entries with smaller size. The end result of all of this is that realized results will start to diverge from the theoretical edge.

Execution risk is, to some extent, a normal cost of doing business. Most executions that are time-sensitive will involve paying the spread on at least one, or maybe both, sides of the trade. In liquid markets, this is a negligible expense unless your strategy trades frequently, but for active traders, even the cost of paying a penny spread in and out can be a significant drag on performance. It is also important to consider what liquidity conditions may exist when you need to get out of a trade quickly. In some thin markets (e.g., illiquid futures, currencies, or small-cap stocks), slippage of 10 percent or more is possible at times. Though this is an extreme example, remember that slippage can be a much larger cost than the novice trader would expect. Another reason that execution may diverge from backtested results is that sometimes, in equities, it may not be possible to locate shares for a short sale, or other regulatory restrictions may interfere with the transaction. It is impossible to fully account for events like this in research and backtesting.

Execution risk can be minimized to some extent through experience and the acquisition of good execution skills. It is also important to invest money in technology, infrastructure, and relationships with the best brokers. A good broker should have competitive rates, but no trader should have to deal with a subpar platform or a poor Internet connection. Furthermore, new ideas in new markets should be explored on very small size to judge the difficulties of execution in the unfamiliar territory. If a trader normally risks $5,000 on a trade, perhaps a foray into a new market would be done in a minuscule size, risking $50 per trade. There may be issues as size is increased, but it is always a good idea to have your mistakes and errors on small size where possible.

Disaster Risk

Some *disaster risks* are fairly common, but some are almost impossible to plan for. Examples of fairly mundane risks are loss of power or Internet connection, a hard drive crash, or some other mechanical failure. It is important to have backup systems and to have a plan that extends several steps further than that. But there are many more serious risks that bear consideration. What would a natural or man-made disaster in a major financial center do to your trading program? What if you could not access your trading accounts for a period of days or weeks? Some of these risks are systemic risks

and, frankly, cannot be managed very well. There is no point in obsessing over things you cannot control, but there is great value in being as prepared as you possibly can be.

Markets Evolve. Do You?

Throughout market history, there have been many ideas that have worked for a while or in one specific environment, and then they stop working, often after being published. Though we try to build our trading programs on principles that will not change, any trading plan is a specific expression of those tendencies expressed in one specific way. It is completely possible that what worked yesterday will not work tomorrow. This is reality. Stuff, as they say, happens. Markets evolve and erase certain kinds of trading edges.

This is why any system, including a discretionary one, should be subject to some sort of control process, similar to the ones used in manufacturing and quality control (see Chapter 12). The idea is to identify normal variation in the system and then to flag events that seem to violate that normal variation. This is not a simple task, as markets are subject to much more randomness than any manufacturing process, and large surprises are normal for many trading systems. However, with a little work and adaptation, it is possible. Monitoring your performance for severe changes or degradation can give early warning and allow time to make some adjustments. If you are trading an idea that stops working, you will incur some losses. This is unavoidable, but a control process will monitor those losses as they evolve and give a cutoff point at which trading should be terminated. The real danger is not recognizing that something has changed, and continuing to pour money into a trading system that has become a bottomless hole.

Event Risk/Tail Risk

Tail risk refers to extreme events in the tails of a distribution. We might observe a commodity trade for several years, and see that most of its daily returns fall within +/−2 percent of the previous day, with the occasional large move out to 4 percent. However, it is entirely possible that that same market could make a 25 percent move tomorrow, out of all proportion to anything in the historical record. Figure 9.4 shows a daily chart of Procter & Gamble (NYSE: PG) covering several months, including the day of the May 6, 2010 Flash Crash. This was a fairly boring blue-chip stock that was not prone to making volatile movements, but nearly an entire decade of price gains were erased in less than 30 minutes, only to be almost completely regained within the next few minutes. This was a move that most market participants probably thought was impossible, and yet it happened. There are other stocks that showed even more dramatic movements on this day, and this type of event happens several times a year in individual stocks.

In many cases, though not all, these large standard deviation events will be driven by geopolitical events or natural disasters and are completely unpredictable. It is the nature of these events that they are complete surprises. In addition, they are also more or less unhedgeable; appropriate hedges do exist for most of these risks, but the costs of applying them continuously will set a hurdle rate that would be difficult to overcome. For

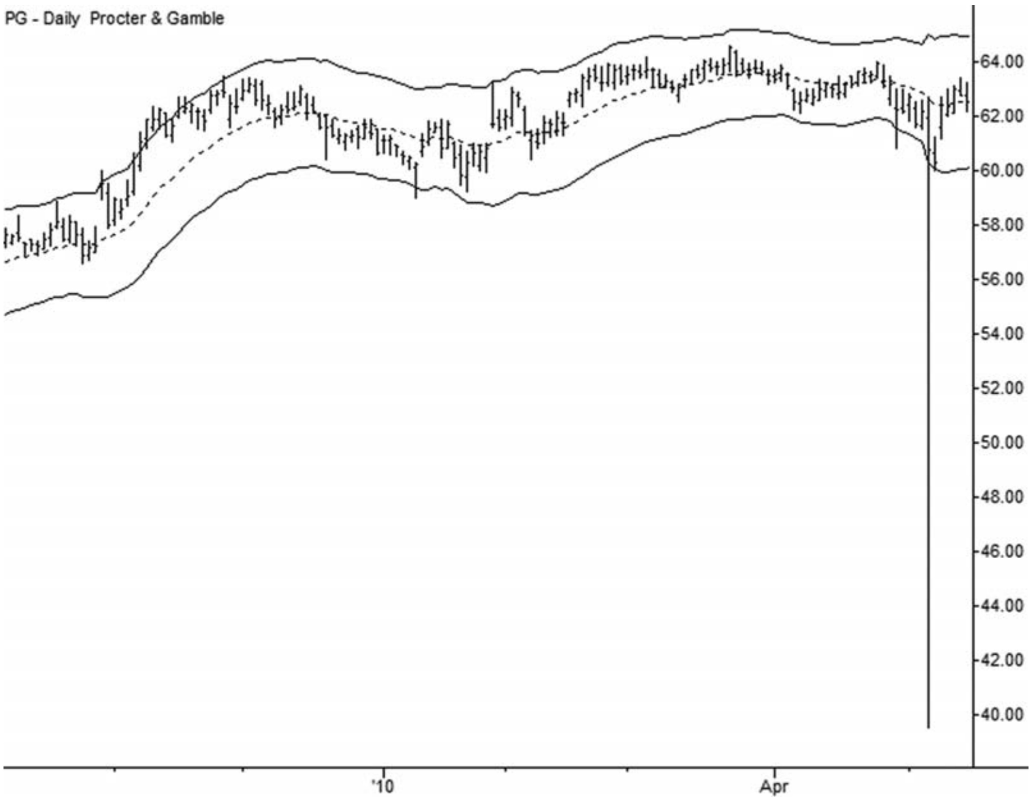

FIGURE 9.4 Daily Chart of PG During the 2010 Flash Crash

instance, the cost of fully hedging a complex portfolio with options can easily run over 10 percent on an annualized basis. With that kind of haircut, you have to ask yourself just how good you think you are. Can you hedge, beat the hedge hurdle, and still make enough to make the venture worth your time? Not likely.

When most traders talk about tail risk, they refer to the kind of events that can put a trader out of business, but remember that tail events are not always bad. Sometimes the trader may be positioned on the right side of the surprise move, and it turns out to be a gift. Though some traders have tried, it is very difficult to structure strategies designed specifically to take advantage of these events, because they are so infrequent and volatile. Even if you are fortunate enough to be on the right side of such a move, there is no guarantee that you will find adequate liquidity to exit at a profit.

Correlation

Correlation is a major concern in portfolio design and analysis, but also may have a larger than expected impact on an active trader holding a handful of positions. Though correlation is usually considered in the context of longer-term portfolios, correlation risk

cannot be ignored by the active trader. For instance, an equity swing trader who intends to hold positions for two to four weeks may often be holding five or more positions, distributed across various sectors and industries. Normally, these positions would be quite likely to each follow its own course, with some small degree of dependence on the overall index. A futures swing trader may hold positions in currencies, metals, and grains that turn out to be much more tightly correlated than expected. If a major event hits the market, they will all move together. In a nutshell, correlation is bad because it destroys diversification effects—diversification is not there exactly when you need it most. A trader may believe she is risking 3 percent of her portfolio on four different positions, but if they are highly correlated, it is much more like one large position risking 12 percent of equity. This is far more than most traders intend to risk on any one position, but they do it all the time when they underestimate correlation and risk, and the risk of *shifting* correlations.

There is a good reason to expect increased correlations in times of financial stress: large pools of money tend to move together in times of crisis. This has been well known and documented for decades, but the rise of ETF products has given more investors access to broad swaths of asset classes. Many of these assets are being pooled in various funds as general risky assets that will be dumped en masse in times of systemic stress. If this trend of aggregation continues, and it seems likely that it will, this effect is likely to intensify. Active traders must always consider the possibility that every position they hold can move against them at the same time.

Liquidity Risk

For most traders, the main danger of low liquidity is that you may be unable to execute at prices you intend. Even moderately sized orders can sometimes move markets far more than expected; if the order is time-sensitive, as in exiting a position at a stop point, heavy slippage can result. This is to be expected if the position represents a significant percentage, or multiple of, an asset's average daily trading volume. What might not be so obvious, and thus is more dangerous, is that lack of liquidity can pose a threat to much smaller positions in normally liquid markets. In addition, this effect has been exacerbated in many markets by the risk of high-frequency trading (HFT) programs that provide what is essentially predatory liquidity. Automated market makers are capable of bidding and offering tight spreads, and then widening those spreads when other orders take liquidity from the order book, at speeds far beyond human reaction time.

Regulatory Risk

Regulatory risk is rarely considered, but can have a catastrophic impact in certain market environments. In more mundane cases, a change to tax laws may impair the effectiveness of a hedge or an offsetting position. Sometimes, an increase in margin requirements for a futures contract may have implications for position sizing and leverage across a portfolio. It is not uncommon to see margin requirements adjusted by the exchanges after a

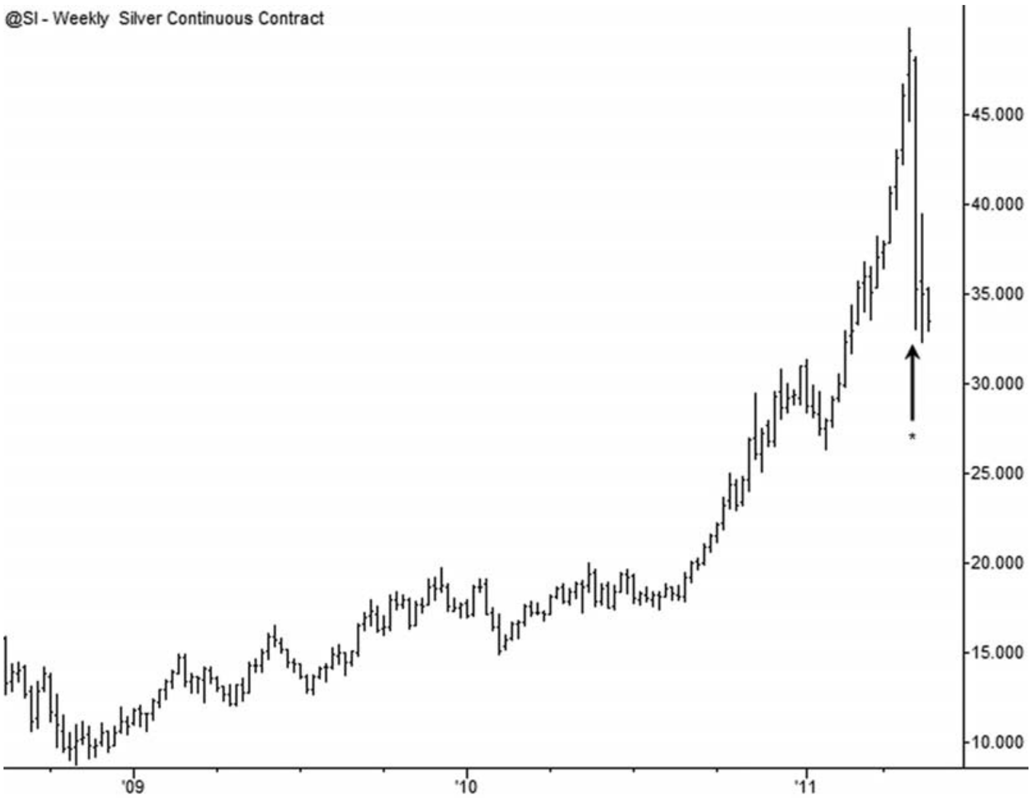

FIGURE 9.5 A Margin Hike in Weekly Silver Futures Breaks an Already Overextended Trend

large price movement, and sometimes these margin changes can break the back of an established trend. Figure 9.5 shows the effect of several margin hikes for Silver futures on the trend in mid-2011. Intuitively, you might expect that such a margin increase would be insignificant, but it can force liquidation and bring large-scale selling pressure into the market. The Hunt brothers' silver fiasco in the early 1980s is one of the best-known examples, but there have been many others as the regulatory environment continually evolves in response to the actions of traders and investors.

In equities, there are frequent restrictions on shorting that may cause actual trading results to diverge from theoretical tests, and there have even been examples where shorting has been banned altogether for a group of stocks or entire regions. In early 2011, some of the cutting edges of regulatory coverage are the issues surrounding HFTs, algorithmic trading programs, and the prevalence of computer-assisted trading programs that seek to capture the liquidity rebates across electronic communication networks (ECNs). It is quite likely that some regulatory move in the next few years will significantly alter the market microstructure and force traders to adapt.

It has also been fairly common to see trades broken in some environments. For instance, there are always traders who buy into large declines like the 2010 Flash Crash,

and then sell their inventory into the rallies. Note to regulators and exchange officials: you want to *encourage* this kind of behavior. When markets are crashing, traders who take the risk of buying and potentially supporting markets should be rewarded for this behavior. Instead, in the Flash Crash, many trades were broken outside of arbitrary ranges that were not made up until the next day. This is capricious and random behavior on the part of the regulators. Imagine, for instance, that you stepped in to buy a stock that was down 40 percent, and you offered most of your inventory out at 30 percent, 20 percent, and 10 percent down. As the market recovered, the stock erased its sell-off, your sell orders were filled, and you went home flat. Now you wake up the next morning and discover that the exchange broke all trades more than, say, 30 percent down. What has happened? Your *buy* order was canceled as though it had never happened, but your sells stand, and you are now *short* the stock from 20 percent under the current price. If you are an individual trader, this kind of loss could put you out of business. This is not an isolated example; regulators and exchange officials make these random decisions several times a quarter in various stocks. There is no incentive for the individual trader to step into the breach and to provide liquidity in the event of severe declines, because your trades may simply be erased the next day.

SUMMARY

Risk management is the first and most important job of any trader. Traders must make clear distinctions between the normal risks associated with any trade and the more extraordinary risks that can potentially destroy a trading account or end a trading career. Many of the psychological struggles traders face come from not really understanding the nature of risk. The trader's job boils down to this: assume the correct kinds of risk at the correct times, and manage those risks appropriately; then, assuming the trader is working within a net positive expectancy framework, profits will accrue.